

Appl. No. 10/036,140
Amdt. dated June 23, 2005
Reply to Office action of April 7, 2005

Amendments to the Specification:

Please replace paragraph [0019] with the following amended paragraph:

[0019] As defined in its standards proposal, SCTP preferably treats problems with failure of the path in the following manner. For example, a count is maintained on the number of retransmissions to a particular destination address without successful acknowledgement. When the count exceeds a configured maximum, the address is declared inactive, notification is given to the application, and the SCTP instance begins to use an alternate address for sending data chunks. Additionally ~~heartbeat~~heartbeat chunks (or heartbeats) are sent periodically to all idle addresses (i.e., alternate addresses defined in the association), and a counter is maintained on the number of ~~heartbeats~~heartbeats sent to an idle destination without receipt of a corresponding heartbeat acknowledgement. When this counter exceeds a configured maximum, that destination address is also declared inactive. Regardless of how or why an address is declared inactive, heartbeats continue to be sent to the inactive destination addresses until an acknowledgement is received, at which point the address can be made active again. The rate of sending heartbeats is tied to the retransmission timeout (RTO) estimation plus an additional delay, to allow heartbeat traffic to be tailored to the needs of the application.

Please replace paragraph [0027] with the following amended paragraph:

[0027] In the event of failure of the node on which the SCTP instance is located, the duplicate instance on the other node in the cluster will take over in a manner which is transparent to the instance on the other end of the association. The other end will only perceive, either on its own through normal processes or alternatively through an express notification from the duplicate instance, that certain of the IP addresses to the receiving instance are no longer functioning (do not provide a ~~heartbeat~~heartbeat), but that one of the alternate addresses will communicate and now becomes the primary address (and hence the primary pathway). In this manner, the multi-homing feature designed for one purpose in the current SCTP protocol is used for a separate purpose by the preferred

Appl. No. 10/036,140
Amdt. dated June 23, 2005
Reply to Office action of April 7, 2005

embodiment of the disclosed system. By allowing the association to be preserved even when the primary node is lost, a level of fault tolerance is provided to the instance itself as well as the fault tolerance designed in to SCTP for the connections making up the association.

Please replace paragraph [0028] with the following amended paragraph:

[0028] Figure 3 illustrates a top-level view of one embodiment of the present disclosure. A first computer 10 and a second computer 10A are both a part of a first cluster ~~[[40]]~~ 110. IP address 12 points to first computer 10 and provides an interface on that node. IP address 14 points to second computer 10A and provides an interface on that node.

Please replace paragraph [0029] with the following amended paragraph:

[0029] It is typical for clustered machines to have a back-door for communication. In a monolithic host there is communication between different physical interfaces across its own bus or internal systems. In a clustered system such as first cluster ~~[[40]]~~ 110, this link includes links (forming an intra-cluster network) between the separate computers, such as intra-cluster connection 16 between first computer 10 and second computer 10A. The intra-cluster connections are cooperatively working together in the cluster. These connections may also be referred to as an intra-cluster communication channel which in the preferred embodiment of the present disclosure provides dynamic update of the state of the association between the two computers. There will be separate instantiations (instances) of SCTP on each computer (10 and 10A) and the state diagram for the associations will be periodically synchronized between the computers in the cluster (by communicating over intra-cluster connection 16).

Please replace paragraph [0030] with the following amended paragraph:

[0030] As with the proposed SCTP standard, the association 30 provides a communication path or multiple communication paths. In SCTP an association may include several possible interfaces. The control software will heartbeat to

**Appl. No. 10/036,140
Amdt. dated June 23, 2005
Reply to Office action of April 7, 2005**

keep track of which interfaces are alive and well. If the primary interface goes down, the system automatically switches to one of the alternate interfaces. In the preferred embodiment of the present disclosure, some of the interfaces are on separate computers. In the illustrated embodiment, not only is the first cluster [[40]] 110 providing a multiple node endpoint, but the second cluster 50 at the other end of association 30 also provides two nodes (computers 20 and 20A having IP addresses 22 and 24). For the purposes of the disclosure the other or opposite endpoint could equally well be a traditional single node on a computer 20 having one or multiple IP addresses (and hence one or multiple interfaces). If the first computer 10 with the primary interface goes down, the opposite endpoint, or more specifically the instance of SCTP on the other end of the association (cluster 50 as illustrated here), will note the failure and send to an alternate interface (the second computer 10A). Because the second computer has a copy of the state information for the association, it is able to step in and receive the information even if its own ability to reach the first computer is compromised. To the endpoint on the second or opposite cluster 50 here, it is transparent as to whether the alternate interface is on the same node or host or on a different node or host. It knows only that the primary interface has failed and information should be sent to the alternate interface, which should have access to the state information to know what to do with the information being sent.